

## ABSTRACT OF THE DISCLOSURE

A method and apparatus for scheduling the transfer of data bursts in a network comprising electronic edge nodes interconnected by bufferless core nodes are disclosed. Each edge node comprises a source node and a sink node, and each core node comprises several  
5 bufferless space switches operated in parallel. Each source node is connected to at least one core node by an upstream link that includes multiple upstream channels. Each core node is connected to at least one sink node by a downstream link that includes multiple downstream channels. Any of the space switches can have either an electronic fabric or a photonic fabric. Each space switch has a master controller, and one of the master controllers in a core node is  
10 designed to function as a core-node controller in addition to its function as a master controller. Each master controller has a burst scheduler operable to compute a schedule for the transfer of data bursts, received from source nodes, to destination sink nodes.

In one mode of operation, each source node determines the bitrate requirements for paths to each sink node and sends bitrate-allocation requests to a core-node controller. A  
15 core-node controller receives requests for bitrate allocations from source nodes and assigns each request to one of the master controllers of the core node. Each master controller then computes burst-transfer permits and sends the permits to corresponding edge nodes. Each permit specifies a burst size, a destination sink node, the time at which the burst should be transmitted from its source node, and the upstream channel to be used. It is emphasized that  
20 in this mode of operation, the source node communicates information on burst streams not on individual bursts.

In another mode of operation, each source node determines a descriptor of each unscheduled waiting data bursts and communicates the burst descriptors of all unscheduled bursts to a controller of a core node. The core-node controller then directs each burst  
25 descriptor to one of the master controllers of the core node which then computes a schedule for burst transfer and sends the schedule to corresponding edge node. The schedule indicates the upstream channel to be used for each burst and the time at which each burst, already waiting at the source node, should be transmitted to the core node.

The scalability of the network is determined, in part, by the processing capacity of the  
30 burst scheduler and techniques for circumventing this limitation are described in this disclosure.